

Robust predictions for an oscillatory bispectrum in Planck 2015 data from transient reductions in the speed of sound of the inflaton

Jesús Torrado^{1,2}, Bin Hu^{3,4} and Ana Achúcarro^{5,6}

¹ *Astronomy Centre, University of Sussex, Brighton BN1 9QH, UK (current address)*

² *Institute for Astronomy, University of Edinburgh, Blackford Hill, Edinburgh EH9 3HJ, UK*

³ *Department of Astronomy, Beijing Normal University, Beijing, 100875, China*

⁴ *Institut de Ciències del Cosmos (ICCUB), Universitat de*

Barcelona (IEEC-UB), Martí i Franquès 1, E08028 Barcelona, Spain

⁵ *Institute Lorentz, Leiden University, PO Box 9506, Leiden 2300 RA, The Netherlands*

⁶ *Dept. of Theoretical Physics, University of the Basque Country UPV-EHU, 48080 Bilbao (Spain)*

We update the search for features in the Cosmic Microwave Background (CMB) power spectrum due to transient reductions in the speed of sound, using Planck 2015 CMB temperature and polarisation data. We enlarge the parameter space to much higher oscillatory frequencies of the feature, and define a robust prior independent of the ansatz for the reduction, guaranteed to reproduce the assumptions of the theoretical model and exhaustive in the regime in which the feature is easily distinguishable from the baseline cosmology. We find a fit to the $\ell \approx 20$ –40 minus/plus structure in Planck TT power spectrum, as well as features spanning along the higher ℓ 's ($\ell \approx 100$ –1500). For the last ones, we compute the correlated features that we expect to find in the CMB bispectrum, and asses their signal-to-noise and correlation to the ISW-lensing secondary bispectrum. We compare our findings to the shape-agnostic oscillatory template tested in Planck 2015, and we comment on some tantalising coincidences with some of the traits described in Planck's 2015 bispectrum data.

PACS numbers: 04.60.-m; 98.80.-k; 98.80.Cq; 98.80.Qc

I. INTRODUCTION

The last data release of Planck [1] has already made public all the data taken by the survey, including polarisation power spectrum data for public analysis, and some results of the analysis of the bispectrum, though not a likelihood for public use. The analyses carried out by the Planck collaboration in the context of primordial fluctuations have not found any strong deviation from the predictions of the canonical single-field slow-roll inflation paradigm. In particular, they found no significant deviation from the vanilla power-law power spectrum [2], neither a detection of any shape of primordial non-Gaussianity [3]. Some hints are reported for small deviations on both data sets, but always in the low signal-to-noise regime, under the significance necessary to claim a detection. Some of those hints persist in Planck 2013 [4, 5] and Planck 2015 (and even WMAP [6]), such as a dip at $\ell \approx 20$ and some small features in the CMB temperature bispectrum, that have been deemed interesting by the Planck collaboration.

Many of the extensions of canonical single-field slow-roll inflation predict [7] correlated features in both the 2- and 3-point correlation functions¹. Notably, in a few cases, the correlations can be computed explicitly [8–12]. When models with correlated features are tested against the data in a joint approach for different observables at the same time, the significance of possible fits is expected to be increase, as has been reported in particular for oscillatory feature searches combining CMB power spectrum and bispectrum [13–15] (see also [16] for a model-independent approach).

This motivates us to update our ongoing search [17–19] for features produced by transient reductions in the speed of sound of the inflaton [9] with the new Planck 2015 temperature and polarisation power spectrum data, in preparation for a joint search including bispectrum data. As a part of it, we have re-evaluated the prior of our search to ensure theoretical self-consistency in a more efficient way (imposed a priori, not a posteriori) and enlarge the parameter space such that it covers all the configurations for which the feature is *distinguishable* from the baseline cosmology. With the results of this updated search, we formulate predictions for the CMB bispectrum that are *robust*, i.e. are guaranteed to be theoretically self consistency and have a very narrow range of oscillatory frequencies.

The present paper is structured as follows: we begin by reviewing the theoretical framework for this family of inflationary features, and describe their shape in the CMB observables (section II A); then, we present our ansatz for the speed of sound reduction (section II B) and discuss the prior that we will employ in our sampling (section II C). After discussing the data sets and methodology with which our search has been conducted (section III), we present and discuss our results for the CMB power spectrum (section IV), and draw from them predictions for the CMB

¹ There is a very extensive literature on this subject, we refer the reader to the recent review [8].

bispectrum (section V) which are discussed in the context of Planck's search for non-Gaussianity (the numerical tools used to carry out CMB bispectrum computation and forecasts are described in appendix A). Finally, we discuss the relevance of our findings and prospects for searches for features of this kind (section VI).

II. THEORETICAL MODEL AND PRIOR

A. Review of the theoretical model

We work in the framework of effective field theory of inflationary perturbations [20], described in terms of the Goldstone boson of time diffeomorphisms, $\pi(t, \mathbf{x})$. This is related to the adiabatic curvature perturbation linearly: $\mathcal{R}(t, \mathbf{x}) = -H(t)\pi(t, \mathbf{x})$, with $H := \dot{a}/a$ and a is the scale factor (from now on, we denote physical time derivatives with an overdot, $\dot{} = d/dt$).

The effective quadratic action for π reads

$$S_2 = M_{\text{Pl}}^2 \int d^4x \frac{\epsilon a^3 H^2}{c_s^2} \left\{ \dot{\pi}^2 - c_s^2 \frac{(\nabla \pi)^2}{a^2} \right\}, \quad (1)$$

where $\epsilon := -\dot{H}/H^2$ and the time-depending speed of sound c_s that appears in the action accounts for the effect of the heavy components of the field space that are made implicit by the effective field theory.

In order to get a physical grasp of the significance of a speed of sound reduction, carrying out explicitly the integration of the heavy mode in a 2-field scenario, one gets

$$c_s = \left(1 + \frac{4\dot{\theta}^2}{M^2 - \dot{\theta}^2} \right)^{-2}, \quad (2)$$

where $\dot{\theta}$ is the angular velocity (the rate of turning) of the background trajectory along the minimum of the potential [21], and M^2 would be the mass squared of the (heavy) modes perpendicular to that trajectory if the trajectory were straight. Thus, soft (adiabatic) *turns* in the inflationary trajectory in field space result in transient reductions of the speed of sound.²

We can re-write the quadratic action (1) as

$$S_2 = S_{2,\text{free}} + M_{\text{Pl}}^2 \int d^4x \epsilon a^3 H^2 (-u \dot{\pi}^2), \quad (3)$$

where $S_{2,\text{free}} := S_2(c_s = 1)$ and we have re-parametrised the varying speed of sound as [24]

$$u := 1 - \frac{1}{c_s^2}, \quad (4)$$

which departs from zero towards *negative* values when the speed of sound departs from unity. Treating the transient speed of sound as a small perturbation of the free action and using the *in-in* formalism [25], one sees that mild changes in the speed of sound seed *features* in the primordial power spectrum of curvature perturbations as [9]

$$\frac{\Delta \mathcal{P}_{\mathcal{R}}}{\mathcal{P}_{\mathcal{R}}}(k) = k \int_{-\infty}^0 d\tau u(\tau) \sin(2k\tau) \quad (5)$$

where $\mathcal{P}_{\mathcal{R}} = H^2/8\pi^2 \epsilon M_{\text{Pl}}^2$ is the featureless nearly-scale-invariant power spectrum corresponding to the constant $c_s = 1$ ($u = 0$) case, $u(\tau)$ departs briefly and softly from zero and back, and τ is conformal time.

One can also write the cubic action for the adiabatic mode:

$$S_3 = M_{\text{Pl}}^2 \int d^4x \epsilon a^3 H^2 \left\{ -2(1-u) s H \pi \dot{\pi}^2 - u \dot{\pi} \left[\dot{\pi}^2 - \frac{(\nabla \pi)^2}{a^2} \right] \right\}, \quad (6)$$

² Sufficiently sharp turns would violate the adiabatic condition that prevents quanta of the heavy degrees of freedom from being produced [22, 23]: $|\dot{c}_s| \ll M|1 - c_s^2|$.

where we have introduced the relative derivative of the speed of sound

$$s := \frac{1}{H} \frac{\dot{c}_s}{c_s}. \quad (7)$$

In the cubic action above, two important assumptions have been made:

- Slow-roll contributions still present in constant-speed-of-sound scenarios are neglected. They come at order $\mathcal{O}(\epsilon^2, \eta^2)$ [26], so in order for this assumption to be correct (i.e. this here being the main contribution to the cubic action), *at least one of* u or s must be significantly larger than the slow-roll parameters, at least at their maximum deviation from zero.
- The cubic action due to the speed of sound reduction is treated perturbatively, so if we want to be sure that higher order terms can be neglected, *both* the speed of sound and its change rate as they appear in the cubic action, u and s , must be significantly smaller than 1 at their maxima.

Summarising:

$$\epsilon, |\eta| \ll \max(|u|_{\max}, |s|_{\max}) \ll 1. \quad (8)$$

In section II C, we discuss how to impose those bounds in a natural way.

It is easy to check that the perturbative limit on $|s|_{\max}$ ensures that the consistency conditions derived in [22, 27, 28] are comfortably satisfied, setting a limit to the sharpness of the reduction at least as stringent as the ones found in those references. Thus, as long as our prior duly imposes those bounds in $|u|_{\max}$ and $|s|_{\max}$, we eliminate the risk of fitting to the data features whose computation can be found a posteriori not to be theoretically consistent.³

From the cubic action above, again using the *in-in* formalism, one can compute the main contribution to the spectrum of the curvature perturbations [9]

$$B_{\mathcal{R}}(k_1, k_2, k_3) = \frac{(2\pi)^4 \mathcal{A}_s^2 M_{\text{Pl}}^6}{(k_1 k_2 k_3)^2} \sum_{i=0}^2 c_i(k_1, k_2, k_3) \left(\frac{k_t}{2}\right)^i \left(\frac{d}{dk_t/2}\right)^i \frac{\Delta \mathcal{P}_{\mathcal{R}}}{\mathcal{P}_{\mathcal{R}}}(k_t/2), \quad (9)$$

where $k_t := k_1 + k_2 + k_3$. The scale-independent *shape coefficients* c_i are:

$$c_0 := -\frac{1}{k_t^2} \left(k_1 k_2 + \text{cyclic}\right) + \frac{1}{4} \frac{1}{k_t} \left(\frac{k_1^3}{k_2 k_3} + \text{cyclic}\right) - \frac{3}{2} \frac{1}{k_t} \left(\frac{k_1 k_2}{k_3} + \text{cyclic}\right) + \frac{1}{4} k_t \left(\frac{1}{k_1} + \text{cyclic}\right) - \frac{5}{4} \quad (10a)$$

$$c_1 := \frac{1}{k_t^2} \left(k_1 k_2 + \text{cyclic}\right) - \frac{19}{32} + \underbrace{\frac{19}{32} - \frac{1}{4} \frac{1}{k_t} \left(\frac{k_2^2 + k_3^2}{k_1} + \text{cyclic}\right)}_{c_{1,\text{sq}}} \quad (10b)$$

$$c_2 := \frac{1}{4} \frac{1}{k_t^2} (k_1^2 + k_2^2 + k_3^2) \quad (10c)$$

where cyclic means the 2 remaining cyclic permutations of the k_i (the missing k_i 's in a term are understood to be implicit, e.g. $k_1 k_2 + \text{cyclic} := k_1 k_2 + k_2 k_3 + k_3 k_1$).

Notice that, unlike in most of the literature, we are not extracting an overall amplitude f_{NL} in front of the bispectrum. We could use $|u|_{\max}$ as a proxy for f_{NL} , redefining $(\Delta \mathcal{P}_{\mathcal{R}}/\mathcal{P}_{\mathcal{R}})^* := 1/|u|_{\max} \Delta \mathcal{P}_{\mathcal{R}}/\mathcal{P}_{\mathcal{R}}$. Also, we notice that there is a non-separable prior on $|u|_{\max}$ and $|s|_{\max}$ determined by eq. (8) (and developed in section II C). This non-separability of the amplitude from the rest of the shape parameters should be taken into account when fitting this template to the data, since the range of amplitudes $|u|_{\max}$ allowed by the prior depends on the value of $|s|_{\max}$ of the tested template (see section II C).

The main result from [9] is thus that *features in the power spectrum and the bispectrum are correlated in a very simple, analytic way*, and that both are easily expressed in terms of a mild, transient reduction of the speed of sound $u(\tau)$ of the adiabatic mode. It is worth remarking that both observables were re-computed in the same theoretical framework using the generalised slow-roll formalism in [18] (and in the case of the power spectrum, also a numerical solution to the equations of motion), and they were found to be consistent with the expressions above, with agreement

³ This is different from the treatment in our previous work [17, 18], and also in [4, 29] for steps in the potential.

improving as the reductions get sharper (large $|s|_{\max}$), i.e. the regime where the generalised slow-roll approximation works best.

Let us discuss a little the appearance of those features in both observables. Let us assume that the speed of sound reduction happens around a particular instant τ_0 , that we will define as the instant of maximum reduction: $u(\tau_0) := -|u|_{\max}$. The rate of change s being limited from below by the slow-roll parameters means that the reduction must be approximately localised around τ_0 . The Fourier transform in eq. (5) turns that localisation into a linear-in- k oscillatory factor $\sin(2k\tau_0)$ for the power spectrum feature, with possibly a small phase if the reduction is not symmetric around τ_0 . The finite span in τ of the reduction imposes a finite envelope on top of those oscillations, the details of which (weight of the tails, symmetry) is determined by the particular shape of $u(\tau)$.

In the bispectrum, all this remains true, the oscillatory factor being $\sin(k_t\tau_0)$. The variation along total scale $k_t = k_1 + k_2 + k_3$ is given mainly by $\frac{\Delta\mathcal{P}_{\mathcal{R}}}{\mathcal{P}_{\mathcal{R}}}$ and its derivatives, so when observed along k_t in a particular direction (i.e. a particular triangular configuration), the feature will look similar to that on the power spectrum: an enveloped oscillation. The amplitude and phase do change across different triangular configurations: the *central* configurations (i.e. those away from the squeezed limit, including the *equilateral* and *folded* limits) are dominated by the term with the second derivative and, if the reduction is not specially sharp (i.e. $|s|_{\max} \sim |u|_{\max}$, instead of $|s|_{\max} \gg |u|_{\max}$), it receives additional small contributions from the rest of the terms, especially from the zeroth derivative; the *squeezed limit* is completely defined by the term $c_{1,\text{sq}}$ in the first derivative alone, which diverges towards that limit as the inverse of the smallest wavelength. Despite there apparently being squeezed contributions from $c_0 \cdot \frac{\Delta\mathcal{P}_{\mathcal{R}}}{\mathcal{P}_{\mathcal{R}}}$, they cancel out, in agreement with the consistency condition [26, 30].

Due to the order of the derivatives, *the oscillations in the squeezed limits are out of phase by $\pi/2$ with respect to those at the central configurations*. This is the main difference with the shapes tested so far on the Planck bispectrum in the 2013 [5] and 2015 [3] data releases (see section V).

All the statements above about the characteristics of the features are independent from the particular ansatz chosen for the reduction, and are illustrated in figure 1. In the next section, we present our case study: a Gaussian reduction of the speed of sound.




FIG. 1: Features in the primordial power spectrum 1(a) and bispectrum 1(b) (with $\frac{(2\pi)^4 \mathcal{A}_s^2 M_{\text{Pl}}^6}{(k_1 k_2 k_3)^2} \mathcal{S} = \mathcal{B}_{\mathcal{R}}$) from a Gaussian reduction in the speed of sound, eq. (11), with parameters $B = -0.024$, $\log \beta = 5.6$ and $\tau_0 = -203$, corresponding to one of our *maxima a posteriori* (see table I). Notice the linear oscillation along the (total) scale for the (bi)spectrum, and the $\pi/2$ phase difference between the squeezed and the central limits of the bispectrum, as discussed at the end of section II.

B. Gaussian ansatz for the reduction

As in our previous work [17–19], we propose a reduction in the speed of sound as a Gaussian in e-folds (or, equivalently, in physical time):⁴

$$u(\tau) = B \exp \{ -\beta(N - N_0)^2 \} = B \exp \left\{ -\beta \left(\log \frac{\tau}{\tau_0} \right)^2 \right\}. \quad (11)$$

This reduction is parametrised by its maximum intensity $B < 0$, a sharpness $\beta > 0$ and an instant of maximum reduction $\tau_0 < 0$ (or, equivalently, N_0). τ_0 fulfils the localisation role explained in the last section. The intensity and the sharpness here are related to the maxima in the reduction $|u|_{\max}$ and its rate of change $|s|_{\max}$ as

$$|u|_{\max} = -B \quad \text{and} \quad |s|_{\max} = \sqrt{\frac{\beta}{2}} \frac{-B}{e^{\frac{1}{2}} - B}. \quad (12)$$

Notice that this functional form has naturally 3 parameters only, (B, β, τ_0) , exactly as many as we used in the last section to characterise a reduction in a model independent way: $(|u|_{\max}, |s|_{\max}, \tau_0)$. Also, a Gaussian is one of the simplest functions that softly departs from zero and returns.

C. Prior

In our previous work [17–19], we imposed a uniform prior directly on the parameters of the Gaussian reduction, and checked that $|s|_{\max} \ll 1$ a posteriori. Since $|s|_{\max}$ depends on both β and B simultaneously, see eq. (12), a rectangular region in (B, β) does not map nicely into one in $(|u|_{\max}, |s|_{\max})$, where the prior motivated by eq. (8) should be imposed. We successfully explored the parameter region of interest, but in an inefficient manner: regions of the parameter space not allowed by the theory were thoroughly explored to later be thrown away.

In this work, we make *tabula rasa* and try to re-approach the prior choice in a model independent way, from the bare consistency requirements of the theoretical framework (see section II A):

$$\epsilon, |\eta| \ll \max(|u|_{\max}, |s|_{\max}) \ll 1. \quad (8)$$

This condition *on-the-maximum* gives the prior a *framing square* shape, see fig. 2(a): above the diagonal $|u|_{\max} = |s|_{\max}$, the limits given by the equation above must be imposed on $|s|_{\max}$, whereas below the diagonal they must be imposed on $|u|_{\max}$.

The lower limit in the expression above depends on our estimation of the slow-roll parameters. The strictly correct treatment of the problem would be imposing a joint prior on $(|u|_{\max}, |s|_{\max}, \epsilon, \eta)$ which would account for the moving lower bound on eq. (8), or, alternatively, an equivalent prior on $(|u|_{\max}, |s|_{\max}, n_s, r)$, since (n_s, r) are the quantities directly determined by the slow-roll parameters.

For the sake of simplicity, in this study we choose to fix that lower bound in eq. (8) to the *central value* of Planck’s estimate for η (the largest of the slow-roll parameters), using temperature and polarisation data and assuming a featureless power spectrum with free running of the spectral index [2] (the most liberal model). The choice of this central value, $\eta = 0.03$, instead of its upper bound as a *much greater than* limit is not necessarily problematic, as long as:

1. The prior density on $\max(|u|_{\max}, |s|_{\max})$ decays fast towards that limit.
2. We avoid feature shapes that may be correlated with the effect of (n_s, r) , so that this estimate of η remains robust.

We shall also keep in mind that values of $\max(|u|_{\max}, |s|_{\max})$ close to the lower bound are associated to the lower part of the confidence level (c.l.) interval for η . The downside of this choice is that by fixing a value for the lower bound that actually depends on the slow-roll parameters we are giving up on the region of the joint prior on $(|u|_{\max}, |s|_{\max}, \epsilon, \eta)$ in which (ϵ, η) are much lower than the boundary and they leave more, unexploited room for $(|u|_{\max}, |s|_{\max})$. This

⁴ Choosing a Gaussian in τ would have been problematic: it would never be exactly zero by τ_0 , as required for the expression of the bispectrum in eq. (9).

should anyway not worry us too much: smaller values for $|u|_{\max}$ and $|s|_{\max}$ make the features harder to detect, and thus the relevant part of the prior would hardly contain any posterior mass.


The last condition above is easily achieved by imposing a minimum sharpness of the reduction and a minimum oscillation frequency, such that the full feature is contained within the observable CMB window and it performs at least a small number of oscillations. In [17], we found this limit for the sharpness in terms of the parameter β for a Gaussian reduction, to be $\log \beta \geq 0$ (see red line in figure 2(a)). The lower limit on the oscillation frequency $|\tau_0| \geq 70$ is discussed at the end of this section.⁵

As for the first condition, we treat it on the same grounds as the upper bound of eq. (8). Both boundaries are *soft* limits, so the prior probability must decrease towards them, instead of falling sharply. Weighting both extremes equally, we would aim to decrease the prior probability by the same amount if we go away from each limit by the same factor, e.g. same probability mass to the left of twice the lower limit, than to the right of half the upper limit. This occurs naturally if the prior is symmetric on the logarithm of the variables, which makes sense, since the lower and the upper limits are two orders of magnitude apart.

A good choice for a prior that has all those ingredients (zero outside an interval, soft fall towards the limits, symmetry) would be a symmetric log-Beta distribution:⁶

$$\max(\log_{10} |u|_{\max}, \log_{10} |s|_{\max}) \sim \text{Beta}(\alpha, \alpha) \quad \text{with} \quad \alpha > 1, \quad \max(|u|_{\max}, |s|_{\max}) \in [0.03, 1]. \quad (13)$$

The lower the value of the shape parameter, $\alpha > 1$, the more disperse the distribution. In this work, we choose $\alpha = 5$, which places the 95% confidence level interval at double/half the boundaries, and the 68% at thrice/third. This fulfils condition 1 above on the robustness of our choice of a lower bound: it leaves Planck's 2- σ upper bound on η ($\eta \approx 0.05$) under the leftmost 5% of prior mass. The density of this prior corresponds to the shading in figure 2(a).



(a) Prior density in the parameters $(\log_{10} |u|_{\max}, \log_{10} |s|_{\max})$. The red line represents the limit $\log \beta = 0$, and the dotted lines mark $\log \beta = 2, 4, 6, \dots$

(b) Region in the parameters $(B, \log \beta)$ sampled with *uniform* prior density in our previous work [17–19]. The new prior density in figure 2(a) is plotted on top (shading), to highlight how much of the old prior is theoretically disfavoured.

FIG. 2: Prior on intensity and sharpness of the speed of sound reduction in this work, 2(a), and in previous works, 2(b).

Let us now compare the new prior with the one we used in our previous work [17–19], which is uniform over the region plotted in figure 2(b). If we plot the density of the new prior on top of the uniform old prior, we see that approximately $2/3$ of its area is unshaded, i.e. has null probability density under the new prior. If we trust that the new prior appropriately accounts for the consistency requirements of the theory, then sampling from the old prior leads to over-sampling theoretically uninteresting regions, while under-sampling the interesting ones. Thus, we consider the present choice more reasonable and efficient, since not only we are more likely to find fits of theoretically allowed features, but we are also able to do it in a fraction of the sampling time.

⁵ We shall not forget that the regime outside $(\log \beta \geq 0, |\tau_0| \geq 70)$ is allowed by the theory, and therefore it must be taken into account in a full evidence computation.

⁶ For a random variable x in the domain $[0, 1]$, $x \sim \text{Beta}(\alpha, \beta)$ has a probability density function $\mathcal{P}(x) = x^{\alpha-1}(1-x)^{\beta-1}/N(\alpha, \beta)$, with $N(\alpha, \beta) := \Gamma(\alpha)\Gamma(\beta)/\Gamma(\alpha + \beta)$.

Finally, let us look at the prior on the instant of maximum reduction $\tau_0 < 0$. The theoretical model imposes no requirements within the range $(-\infty, 0)$. Still it makes sense to restrict ourselves to a finite range from the search point of view:

- We impose a minimum number of oscillations along the range of scales visible in the CMB power spectrum, $\Delta k \approx 0.2 \text{ Mpc}^{-1}$, in order to avoid that the features mimic other effects of the background cosmology, and so to limit the degeneracies between the feature parameters and the background cosmology in the posterior. In particular, we choose that the feature performs 4 full oscillations, which sets an upper bound $\tau_0 < -70$.⁷
- If the reduction happens too early, it will not leave any trace on the observable scales of CMB window. One can easily check that, for reasonable values of $|u|_{\text{max}}$ and $|s|_{\text{max}}$, features happening before $\tau_0 = -8000$ leave no trace in the CMB.

Regarding the density of the prior on τ_0 , the two natural choices would be either a uniform prior on τ_0 or a uniform prior on $\log(-\tau_0)$. From a Bayesian point of view, a uniform prior in τ_0 means no preference on the instant of maximum reduction *in conformal time*, whereas, since $t \propto N \propto \log(-\tau)$, a uniform prior on $\log(-\tau_0)$ means no preference *in physical time* (or in e-folds). From a physical point of view, both make sense, depending on which of t or τ one considers the natural time scale of inflation. The interesting range of τ_0 covers 3 orders of magnitude, so the balance may be tilted towards a log-uniform choice, but we sample from both cases to keep our analysis robust. Also, in our previous work [17–19] and also here (see section IV), we find that τ_0 is well constrained by the data, so the choice between the priors is not a vital one.

III. DATA SETS AND SAMPLING METHODOLOGY FOR THE POWER SPECTRUM

The features from the reduction are computed using a *fast Fourier transform* to perform the integral of the reduction in eq. (5). The primordial power spectrum is then fed to a modified version of the CAMB Boltzmann code [31, 32] (<http://camb.info>), in which we implemented an adaptive precision aimed at making the computation of the CMB power spectra faster by only sampling the k 's where necessary, and as dense as necessary, and the same for ℓ .

The features are fitted to the *unbinned* CMB TT, TE and EE power spectra of the Planck 2015 data release [1, 33]. The inclusion of the polarised spectra is an update on the previous searches that we performed using Planck's 2013 data [17, 18]. The use of the unbinned likelihoods is justified by the high oscillatory frequency that the features can reach: the $\Delta\ell = 30$ binning of the multipoles corresponds roughly to a binning of $\Delta k = 2 \times 10^{-3} \text{ Mpc}^{-1}$ in primordial scales, which is smaller than a full oscillation as soon as $\tau_0 > 1500$, and we do explore much higher values.

The sampling is performed with the sampler/integrator **PolyChord** [34], which was chosen specially because of its multi-modal sampling capabilities, since we know the likelihood to be multimodal from previous searches [17, 18]. Handling of the theory and likelihood codes and the sampler is performed with **CosmoChord**, a modified version of **CosmoMC** [35] that incorporates **PolyChord** as a sampler.⁸

For the sake of performance, we fix the value of the nuisance parameters of the Planck 2015 likelihood to their best fit achieved by the Planck Monte Carlo sample with binned, polarised baseline likelihood (**lowTEB + plikHM_TTEEE**) and baseline ΛCDM model.^{9 10} When not sampled (e.g. in the bispectrum study), the cosmological parameters are fixed to the best fit of that same sample.

IV. RESULTS OF FITS TO THE POWER SPECTRUM

We have performed the sampling on the CMB power spectrum data as described in the last section, varying the baseline ΛCDM cosmological parameters ($\Omega_b h^2$, $\Omega_c h^2$, θ_{MC} , τ_{reio} , $\log A_s$, n_s) over a wide uniform prior, and the feature parameters (τ_0 , $|u|_{\text{max}}$, $|s|_{\text{max}}$) using the prior described in section IIC for the last two, with the additional

⁷ This reinforces our upper limits on the slow-roll parameters discussed above, because in particular this way we avoid degeneracies with n_s and r .

⁸ Our last search [19] was conducted with the **MultiNest** nested sampling algorithm [36–38]. The **PolyChord** sampler used in this work is an improvement on **MultiNest**, that it is tailored for high-dimensional parameter spaces, thanks to the use of slice sampling at each iteration to sample within the hard likelihood constraint of nested sampling.

⁹ See table 2.6 in https://wiki.cosmos.esa.int/planckpla2015/images/f/f7/Baseline_params_table_2015_limit68.pdf.

¹⁰ We have verified that varying the nuisance parameters has a negligible effect on our results.

Mode name	$ \tau_0 $	$\log_{10} u _{\max}$	$\log_{10} s _{\max}$	$-10^{-2} B$	$\log \beta$	$\Delta\chi^2_{\text{MAP}}$
100	[100 (102) 105]	[-1.82 (-1.59) - 1.47]	[-1.11 (-0.91) - 0.62]	[1.5 (2.6) 3.4]	[4.4 (4.8) 6.8]	-22
200	[195 (203) 207]	[-2.16 (-1.62) - 1.44]	[-1.01 (-0.78) - 0.57]	[0.7 (2.4) 3.7]	[4.5 (5.6) 8.2]	-16
800	[770 (801) 830]	[-2.06 (-1.37) - 0.77]	[-0.94 (-0.53) - 0.47]	[0.1 (4.3) 17.0]	[2.0 (5.6) 8.2]	-12
1000	[935 (1099) 1631]	[-2.78 (-0.51) - 0.45]	[-1.03 (-0.63) - 0.54]	[0.1 (31) 35]	[0 (1.5) 7.2]	-9

TABLE I: 68% confidence level intervals and *maxima a posteriori* (MAP, in parenthesis) for the modes described in the text and visible in figure 3. The $\Delta\chi^2$ of the MAP's are given with respect to the best fit of the baseline model to a featureless power spectrum. We also provide c.l. intervals for the derived Gaussian ansatz parameters $(B, \log \beta)$. The c.l. intervals of τ_0 for modes 100 and 200 correspond to a Gaussian in $\log |\tau_0|$. Notice the similarity of the bounds on $|s|_{\max}$ along the table: they all correspond approximately to the prior limits (sec. II C).

constraint $\log \beta \leq 14$, imposed for computational reasons but with negligible effect on the posterior modes (it only affects features with an amplitude $< 10^{-3}$).

As stated in section II C, we have sampled *twice*, with two different priors for τ_0 : one is a uniform prior on $|\tau_0|$, which assigns equal probability for a reduction occurring at any *conformal* time, and another with a uniform prior in $\log_{10} |\tau_0|$, which assigns equal probability for a reduction occurring at any *physical* time. Both cases are physically well-motivated. The result of both samples can be seen in figure 3, and the most relevant modes are shown in table I. The reference value $\chi^2 = 69311$ for the effective χ^2 , used in figure 3 and table I, has been obtained from a run with the same likelihood and a featureless primordial power spectrum. Those differences in χ^2 are shown as an approximate reference, since we have not used a thorough maximisation algorithm. We did not found a clearly larger-than-one Bayesian ratio, so from the model selection point of view this model is still disfavoured when considering power spectrum data only.

We found no significant degeneracies between the parameters of the feature and those of the baseline cosmological model; the correlation coefficients stay below $|\rho| < 0.1$ for most combinations, and only for some combinations with $(\Omega_b h^2, \Omega_c h^2, n_s)$ the correlation coefficient grows up to $|\rho| \leq 0.18$, which is still smaller than what was found in fits to the 2013 data [18]. This reinforces the assumptions made in section II C to avoid those degeneracies, namely the lower bounds $\tau_0 \geq 70$ and $\log \beta \geq 0$, which together enforce a minimum number of oscillations to occur within the CMB window.

Looking at the marginalised posterior for $|\tau_0|$, we identify the following modes (see table I):


Low $|\tau_0|$: two modes at $\tau_0 \sim -100, -200$. They both have a very well-determined oscillation frequency τ_0 and an amplitude $|u|_{\max}$ of a few 0.01's. Due to their low τ_0 , we take their confidence level intervals from the log-uniform- τ_0 sample, where they are better resolved. Both modes correspond to the sharp regime, $|s|_{\max} \gg |u|_{\max}$, or high β .

High $|\tau_0|$: one mode at $\tau_0 \sim -800$, with characteristics similar to the two modes above, but worse χ^2 and looser constraints on the parameters. Also, a much broader mode with $\tau_0 \sim -1100$, with a wide posterior on τ_0 , unbounded amplitude (constrained by the prior), and a clearly lower sharpness β than the rest of the modes, which places it in the not-so-sharp regime, $|u|_{\max} \approx |s|_{\max}$. Despite their different regime, the boundary between these two modes is not clearly defined, so we have imposed it at $\tau_0 = -840$ – thus their 68% c.l. limits on τ_0 are just an approximation.


On the very high $|\tau_0| > 2000$ region, we do not find any significant mode. This is probably due to their high oscillatory frequency: the transfer functions are almost constant with respect to them, so their projection on the CMB sky smears out most of their intensity, needing too high values of $|u|_{\max} = |B|$, disfavoured by the prior.

In all the modes above, $|s|_{\max}$ is constrained by the prior only. This lack of predictivity on $|s|_{\max}$ was already observed in our previous work with Planck 2013 data [17–19]; there, it appeared as a degeneracy between the parameters $(B, \log \beta)$ of the Gaussian ansatz of eq. (11). Moving along that degeneracy eventually saturated the $|s|_{\max} < 1$ bound, which is avoided now by the new and more realistic prior. That degeneracy still persists, in a milder version, between $\log_{10} |u|_{\max}$ and $\log_{10} |s|_{\max}$. As explained in [17, 18], the degeneracy was caused by the fact that a simultaneous increase in $|B|$ and $\log \beta$ produces almost no changes in the aspect of the feature in the CMB power spectrum [18, fig. 9]: a larger $\log \beta$ shifts the mode towards smaller scales, where damping and lensing erases most of the primordial information, while a larger value of $|B|$ keeps the power at larger scales constant. The death of that effect at hands of the new prior is clearly illustrated by the difference between the current $(\tau_0, \log \beta)$ profile in figure 3 and the corresponding ones in our previous work: [17, fig. 1], [18, fig. 5] and [19, fig. 2] – in the last ones, the mode continues towards higher values of $\log \beta$ with almost constant χ^2 , well past the $|s|_{\max} = 1$ mark.

Comparing this results with our previous searches in Planck 2013 and WiggleZ data [17–19], we see that the modes



(a)



(b)

FIG. 3: Marginalised 1d posteriors and 2d χ^2 scatter plots for the feature parameters $(\tau_0, |u|_{\max}, |s|_{\max})$ and the derived parameter $\log \beta$ of eq. (11). The prior on $(|u|_{\max}, |s|_{\max})$ is described in section II C, with the 1d marginalised prior in dashed red line – compare the $(|u|_{\max}, |s|_{\max})$ posterior with the prior in figure 2, and notice how almost nothing is learned on $|s|_{\max}$ from the data. The prior on τ_0 is either log-uniform, 3(a), or uniform, 3(b). The colour scale shows the $\Delta\chi^2$ of the unbinned, polarised Planck 2015 likelihood, and differences are given with respect to the best fit of the baseline model to a featureless power spectrum. The modes observed along τ_0 are described in table I.

at $\tau_0 \approx -100, 200$ correspond respectively to the modes \mathcal{E}, \mathcal{C} already found there.¹¹ Mode \mathcal{A} appears as a very faint mode with $\tau_0 \approx -377$ and $|u|_{\max} \approx 0.02$. However, modes \mathcal{B} and \mathcal{D} of Planck 2013 have no corresponding significant signal in Planck 2015, neither does the mode found at $\tau_0 \approx -540$ in [19]. To check whether those modes are still present in Planck 2015 but have been suppressed by the new prior, we re-ran the chains with the old non-realistic prior, uniform on $(B, \log \beta)$, and the binned likelihoods – we still found no trace of 2013’s modes \mathcal{B} or \mathcal{D} , but we did find a mode close to $\tau_0 \approx -540$, albeit with a very high $|s|_{\max}$ that would discard it under the new prior. The disappearance of mode \mathcal{B} may be related to that mode’s benefiting from the spurious wiggle at $\ell \sim 1800$ in Planck 2013’s TT power spectrum.

To assess the effect of the new high- ℓ CMB polarisation data in our samples, we repeated the analysis of the uniform- τ_0 case with Planck 2015’s unbinned TT power spectrum likelihood plus the low- ℓ polarised likelihood. We found that the polarised data enhances the mode 100 while it significantly dampens the mode 1000, which shows up more intensely and with a sharper τ_0 c.l. interval when using TT+lowTEB. The other two modes do not receive a large correction.




FIG. 4: Differences between the best fit of the Λ CDM baseline model to Planck 2015 power spectrum (polarised low- and high- ℓ likelihoods [39]) and the MAP’s of modes 100 (blue), 200 (orange), 800 (green) and 1000 (red) from table I. Notice how mode 1000 (red) fits the minus/plus structure at $\ell \approx 20$ –40, how mode 800 fits some apparent wiggles at the first acoustic peak, and how modes 100 and 200 fit small deviations from the baseline model across a higher range of multipoles, especially the minus/plus structure at $\ell \approx 800$ –900 in TT, a similar structure at $\ell \approx 350$ –450 in EE.

The residuals of these modes with respect to the best fit of a featureless Λ CDM baseline model are shown in figure 4. We can see that modes 100 and 200 span across most of the multipole range, their main coincidence being fitting the minus/plus structure in TT at $\ell \approx 800$ –900 and in EE at $\ell \approx 350$ –450. The mode 800 is restricted to the first acoustic peak and fits a small number of apparent wiggles in the data. The *maximum a posteriori* (MAP) of 1000 tries to fit the dip in temperature at $\ell \approx 20$ and the following peak at $\ell \approx 40$, at the cost of raising the power at $\ell \approx 10$ and below (a similar feature in 2015 data has been reported in [2, 40, 41]). Despite the goodness of the fit,

¹¹ Notice that mode \mathcal{E} on Planck 2013 was previously discarded due to its low $|u|_{\max}$ and $|s|_{\max}$. The corresponding one in Planck 2015 does not have that problem, since at least $|s|_{\max}$ is large enough.

this is in conflict with the apparent lack of power at very small multipoles seen in Planck’s data, and may impose an even more stringent upper limit on the relative amplitude r of the tensor primordial power spectrum. We leave the study of this possibility for future work.

We could ask whether two or more of those modes could be present in the data simultaneously. This would correspond to the case of the inflaton suffering two consecutive reductions in its sound speed, e.g. due to two consecutive turns in field space. The complete answer to that question would come from a fit of two simultaneous features with the restriction that they do not overlap in τ . Their respective features in the CMB power spectrum may or may not overlap. The subset of the parameter space in which the power spectrum features do not overlap can be characterised using the present search: any pair of the modes that we found that do not overlap either on τ or on the power spectrum could have occurred together. Looking at figures 4 and 5, we observe that the only two possible combinations would be those of mode *1000* with either *100* or *200*. Together, they would greatly improve the fit to the data by $\chi^2 \sim -30$, but at the cost of introducing 6 new parameters in the cosmological model.




FIG. 5: Speed of sound reduction in terms of $u = 1 - c_s^{-2}$ for the modes described in table I, in logarithmic scale for τ , and with the same colours as in figure 4 (the correspondence between colours and τ_0 ’s is here obvious).

V. PREDICTED BISPECTRUM FEATURES

We have computed the CMB temperature bispectrum (TTT) (see e.g. figure 6) using an extension of the expansion in total scale proposed in [42], described in appendix A. As expected, and similarly to what happens for the power spectrum, we find the CMB TTT bispectrum to be close to the primordial oscillatory shape described in sec. II and fig. 1(b), modulated by the transfer functions.

Due to the the lack of a publicly released bispectrum likelihood, we have not been able to perform a joint analysis of the power spectrum and bispectrum. But already at this point, we can use the posterior modes in the last section (see table I) to make predictions for future searches in the bispectrum and to compare them to present searches of similar templates, if any. The basis of those predictions are the narrow constraints on the oscillatory frequencies $|\tau_0|$ of the features, and their rather well-defined intensity $|u|_{\max}$, specially for modes *100* and *200*, but also for mode *800* to a lesser extent. If a fit to the bispectrum of this kind of features hits any of these thin regions in τ_0 and shows a similar intensity $|u|_{\max}$, this would strongly hint towards the presence of a reduction in c_s in the regime considered here. As we stated in the previous section, the power spectrum data is not able to put constrain $|s|_{\max}$ beyond its prior. Thus, we cannot predict a more concrete value for it.

We would like to specially remark modes *100* and *200* (see table I) as predictions for a signal in the bispectrum. Their TTT bispectra (see fig. 6) approximately presents some of the characteristics described in the reconstructed Planck bispectrum (sec. 6.2.1 in [3, sec. 6.2.1]: a plus-minus structure in the equilateral limit at the ℓ ’s corresponding to the first acoustic peak, $\ell_{\text{total}} \in [400, 1200]$, and a negative peak (though preceded by a positive one) associated to the equilateral third acoustic peak, $\ell_{\text{total}} \in [2300, 3000]$. They also present additional structure in other limits and scales, where nothing has been particularly reported by the Planck collaboration in [3, 5], except for a mention of small features in the folded limit and the squeezed limit, the last one claimed to be associated to the ISW-lensing




FIG. 6: Equilateral, squeezed and folded limits of the CMB TTT bispectrum of the *maximum a posteriori* of modes 100 (blue) and 200 (orange). The x axis is the total scale $\ell_t = \ell_1 + \ell_2 + \ell_3$ and the bispectra are weighted by a constant shape in the Sachs-Wolfe approximation [43, 44]. In the equilateral limit, notice how both bispectra present a plus-minus structure in the first acoustic peak ($\ell \in [400, 1200]$) and a negative peak at $\ell_t \approx 2600$, similar to what is described on Planck TTT bispectrum data [3].

Mode	S/N	S/N _{squeezed}	Corr cos	Corr sin	Corr ISW-l
100	7.4	4.5	-0.26	-0.59	-0.03
200	7.5	4.0	-0.21	-0.65	-0.04

TABLE II: Signal-to-noise in the TTT bispectrum of the *maxima a posteriori* (MAP) of modes 100 and 200 and their isolated respective squeezed contributions. The signal-to-noise is referred to the amplitude $|u|_{\max}$ of the MAP for each mode, i.e. a signal-to-noise of 5 for an amplitude $|u|_{\max} = 0.1$ would mean an error bar of $0.1/5 = 0.02$ on that amplitude. We also show the correlation on the TTT bispectrum between the MAP of each mode with the *constant feature* tested by Planck 2015 [3, eq. (15)], with phases corresponding to the cosine and sine cases, as well as the correlation with the ISW-lensing secondary bispectrum.

secondary bispectrum. The coincidence between the bispectra predicted by modes 100 and 200 and the description of Planck 2015’s bispectrum is an interesting result, given that the predicted features come from a fit to the power spectrum only.

Using an idealised version of Planck’s effective beam and noise, we can assess the likelihood that our predictions are found, as well as their correlations with other primordial and secondary templates that have already been searched for. To do that, we use the Fisher matrix formalism (for details on these calculations, which only include the TTT bispectrum, see appendix A 5). The results for modes 100 and 200 are shown in table II. The signal to noise for both modes is approximately the same, and it could possibly grant a detection if this template was tested against data containing the corresponding physical signal. We can also observe that between one half and two thirds of the signal-to-noise stems from the divergent squeezed limit only. We have also computed the correlation between modes 100 and 200 and the ISW-lensing bispectrum, and found it to be very small despite the oscillatory nature in the squeezed limit of both the ISW-lensing bispectrum and our template.

Direct searches for features have been performed in the bispectrum of both data releases of Planck, first using only the TTT bispectrum [5] and later including polarisation and much higher oscillatory frequencies [3]. There, it was stated that oscillatory features that connected the aforementioned structure found in Planck’s bispectrum achieved higher significance, but in neither of those cases a fit was found with a significance high enough to be called a detection; nevertheless the results from fits of oscillatory features were deemed “interesting hints of non-Gaussianity”.

We can speculate whether our predictions are consistent with those hints. In particular, let us look at the linearly-oscillating templates tested by them, whose frequency very precisely satisfies $\omega \approx |\tau_0|$. None of the templates tested by Planck present the shape weighting and difference in phase between limits of our shapes, or their particular enveloping, that enhances the signal at high- ℓ . So we focus on the simplest case of a *constant feature* [3, eq. (15)] $B(k_1, k_2, k_3) \propto \sin[\omega(k_1 + k_2 + k_3) + \phi]/(k_1 k_2 k_3)^2$, unenveloped and shape-agnostic. Its correlations with our modes 100 and 200 for a frequency $\omega \approx |\tau_0|$ and the sine and cosine phases are shown in table II. As we can see, the sine case is the most highly (anti) correlated one.¹²

Interestingly, the Planck collaboration do find a peak at $\omega \approx 100$, with a phase close to zero (especially in polarisation; the phase is not so small in temperature) and a negative amplitude with signal to noise of the expected order of magnitude (~ 0.5 times the signal to noise of our templates). We find this coincidence tantalising, and look forward to testing our templates against Planck data directly in a joint search.

VI. CONCLUSIONS AND DISCUSSION

We have updated our ongoing search for features from transient reductions in the speed of sound of the inflaton, with the new Planck 2015 polarised power spectrum data. We have proposed and explored a prior that exhausts the regime in which the feature would be clearly distinguishable from the baseline cosmology. Since the prior is exhaustive and Planck’s temperature power spectrum is cosmic-variance-limited for almost all the range that is relevant for inflationary features, we can consider these results definitive, at least until higher signal-to-noise polarisation data is available for multipoles in the range $\ell = 500\text{--}1500$.

We have found some modes that have a very well constrained oscillation frequency and a rather well-defined amplitude, whereas their sharpness, in terms of $|s|_{\max}$, is not constrained by the data but by the prior, which comes from the theoretical self-consistency. The predicted correlated bispectra of two of these modes show traits similar to those described in Planck’s TTT bispectrum; in addition, Planck’s search for linearly-oscillatory features picks up the frequency, sign and approximate phase of one of them.

This promising result motivates us to repeat the present search in the future, including Planck’s temperature and polarisation bispectra, and using the prior described in section II C. Such a search should also expand to regions of higher $|\tau_0|$ (higher oscillatory frequency) where nothing was found in the power spectrum – the amplitude of the bispectrum features, contrary to that of the power spectrum’s, is proportional to the oscillatory frequency due to the derivatives in eq. (9) [12], and this significantly enhances the signal-to-noise of highly oscillatory features [42]. If the features corresponding to these modes are actually present in the data, combined searches in both the power spectrum and bispectrum are expected to greatly raise the significance of the fits [13, 15], hopefully to clearly detection-like levels.

If that combined search still fails to deliver enough significance, we will have to wait until larger tomographic data sets are available, such as 21 cm tomography [45, 46] or the next generation of Large Scale Structure surveys [47, 48].

Acknowledgments

We thank Vicente Atal, Antony Lewis, Moritz Münchmeyer, Pablo Ortiz, Gonzalo Palma, Julien Peloton and Donough Regan for various helpful discussions. JT and BH thank the Institute Lorentz of the University of Leiden for its hospitality. JT acknowledges support from the European Research Council under the European Union’s Seventh Framework Programme (FP/2007-2013) / ERC Grant Agreement No. [616170], and the Engineering and Physical Sciences Research Council [grant number EP/I036575/1]. BH is partially supported by the Chinese National Youth Thousand Talents Program and the Spanish Programa Beatriz de Pinós. AA acknowledges support by the Spanish Ministerio de Economía y Competitividad (FPA2015-64041-C2-1-P) and Consolider Ingenio Programme EPI (CSD2010-00064). AA’s research was supported by a grant from the Simons Foundation, by the Foundation for Fundamental Research on Matter (FOM) and the Netherlands Organization for Scientific Research (NWO/OCW).

¹² This is apparently at odds with the fact that most of our signal-to-noise comes from the squeezed limit, which would correspond to the cosine case; but not so surprising when one considers that, on top of the different enveloping, there is a strong difference in scaling towards the squeezed limit between our template (divergent at $\min_i(k_i) \rightarrow 0$) and the constant feature (constant).

Appendix A: Bispectrum computation

1. Review of expansion in total scale

We attempt to apply the method proposed in [42], based on an expansion in the total scale $k_t := k_1 + k_2 + k_3$. There, one assumes that the primordial bispectrum can be written such that the shape function does depend on the total scale only, i.e.

$$B_{\mathcal{R}}(k_1, k_2, k_3) = \frac{(2\pi)^4 \mathcal{A}_s^2}{(k_1 k_2 k_3)^2} S(k_t). \quad (\text{A1})$$

Then one expands the shape function in a Fourier series in an interval $[k_{t,\min}, k_{t,\max}]$ in whose extremes the shape function is zero, up to a sufficient order n_{\max} :

$$S(k_t) = \sum_{n=0}^{n_{\max}} [\alpha_n c_n(k_t) + \beta_n s_n(k_t)], \quad (\text{A2})$$

where we have abbreviated

$$c_n(k) := \cos\left(2\pi n \frac{k}{k_{t,\max} - k_{t,\min}}\right), \quad s_n(k) := \sin\left(2\pi n \frac{k}{k_{t,\max} - k_{t,\min}}\right). \quad (\text{A3})$$

The coefficients of the Fourier series are

$$\alpha_n = \frac{2}{k_{t,\max} - k_{t,\min}} \int_{k_{t,\min}}^{k_{t,\max}} dk_t S(k_t) c_n(k_t), \quad (\alpha_n \rightarrow \beta_n, c_n(k_t) \rightarrow s_n(k_t)). \quad (\text{A4})$$

A crucial advantage of this method is that the sine and cosine in the total scale are separable:

$$c_n(k_t) = c_n(k_1) c_n(k_2) c_n(k_3) - [s_n(k_1) s_n(k_2) c_n(k_3) + \text{cyclic}] \quad (\text{A5})$$

$$s_n(k_t) = -s_n(k_1) s_n(k_2) s_n(k_3) + [c_n(k_1) c_n(k_2) s_n(k_3) + \text{cyclic}], \quad (\text{A6})$$

where cyclic means the 2 remaining cyclic permutations of the k_i .

Now, remember that the primordial bispectrum gets projected to the reduced CMB bispectrum as

$$b_{\ell_1 \ell_2 \ell_3} = \left(\frac{2}{\pi}\right)^3 \int dr r^2 \int dk_1 dk_2 dk_3 (k_1 k_2 k_3)^2 B_{\mathcal{R}}(k_1, k_2, k_3) \prod_{i=1}^3 \Delta_{\ell_i}(k_i) j_{\ell_i}(k_i r). \quad (\text{A7})$$

Defining

$$\mathcal{C}_{\ell n} := \frac{2}{\pi} \int dk j_{\ell}(kr) \Delta_{\ell}(k) c_n(k), \quad (\mathcal{C}_{\ell n} \rightarrow \mathcal{S}_{\ell n}, c_n(k) \rightarrow s_n(k)). \quad (\text{A8})$$

and, equivalently

$$\mathcal{C}_{\ell_1 \ell_2 \ell_3, n} := (2\pi)^4 \int dr r^2 \left[\mathcal{C}_{\ell_1 n} \mathcal{C}_{\ell_2 n} \mathcal{C}_{\ell_3 n} - (\mathcal{S}_{\ell_1 n} \mathcal{S}_{\ell_2 n} \mathcal{C}_{\ell_3 n} + \text{cyclic}) \right] \quad (\text{A9})$$

$$\mathcal{S}_{\ell_1 \ell_2 \ell_3, n} := (2\pi)^4 \int dr r^2 \left[-\mathcal{S}_{\ell_1 n} \mathcal{S}_{\ell_2 n} \mathcal{S}_{\ell_3 n} + (\mathcal{C}_{\ell_1 n} \mathcal{C}_{\ell_2 n} \mathcal{S}_{\ell_3 n} + \text{cyclic}) \right], \quad (\text{A10})$$

where cyclic means the 2 remaining cyclic permutations of the ℓ_i . The final reduced bispectrum is

$$b_{\ell_1 \ell_2 \ell_3} = \mathcal{A}_s^2 \sum_{n=0}^{n_{\max}} (\alpha_n \mathcal{C}_{\ell_1 \ell_2 \ell_3, n} + \beta_n \mathcal{S}_{\ell_1 \ell_2 \ell_3, n}). \quad (\text{A11})$$

The reduced bispectrum is thus *separable*, but there is an additional advantage: whatever model parameters the primordial shape depends upon are now contained in the Fourier coefficients α_n and β_n (and, indirectly, in the choice of n_{\max} and the interval $[k_{t,\min}, k_{t,\max}]$). Thus, if we want to compute the CMB bispectrum for different values of the primordial model parameters, while keeping the background cosmology unchanged, we only need to re-calculate the Fourier coefficients, and we can re-use already pre-computed and stored, projected Fourier modes $\mathcal{C}_{\ell_1 \ell_2 \ell_3, n}$ and $\mathcal{S}_{\ell_1 \ell_2 \ell_3, n}$.

2. Extension and applicability to our bispectrum template

Let us now write a slightly more complicated template:

$$B_{\mathcal{R}}(k_1, k_2, k_3) = \frac{(2\pi)^4 \mathcal{A}_s^2}{(k_1 k_2 k_3)^2} [f(k_1)g(k_2)h(k_3) + \text{perms}] S(k_t), \quad (\text{A12})$$

where perms here runs over the possible combinations of the three functions and the three momenta. For symmetry reasons, this is the way a separable factor would take; e.g. the simplest case would be $k_1 + k_2 + k_3$, which corresponds to $f, g, h = k, 1, 1$, and $k_1 k_2 k_3$ would correspond to $f = g = h = 3^{-1/3} k$.

Now let's promote the $\mathcal{C}_{\ell n}$ and $\mathcal{S}_{\ell n}$ to operators over functions of a single k_i :

$$\mathcal{C}_{\ell n}[f] := \frac{2}{\pi} \int dk j_{\ell}(kr) \Delta_{\ell}(k) c_n(k) f(k), \quad (\mathcal{C}_{\ell n}[f] \rightarrow \mathcal{S}_{\ell n}[f], c_n(k) \rightarrow s_n(k)). \quad (\text{A13})$$

and, equivalently

$$\mathcal{C}_{\ell_1 \ell_2 \ell_3, n}[f, g, h] := (2\pi)^4 \int dr r^2 \sum_{(f, g, h)} \left[\mathcal{C}_{\ell_1 n}[f] \mathcal{C}_{\ell_2 n}[g] \mathcal{C}_{\ell_3 n}[h] - \left(\mathcal{S}_{\ell_1 n}[f] \mathcal{S}_{\ell_2 n}[g] \mathcal{C}_{\ell_3 n}[h] + \text{cyclic} \right) \right] \quad (\text{A14})$$

$$\mathcal{S}_{\ell_1 \ell_2 \ell_3, n}[f, g, h] := (2\pi)^4 \int dr r^2 \sum_{(f, g, h)} \left[-\mathcal{S}_{\ell_1 n}[f] \mathcal{S}_{\ell_2 n}[g] \mathcal{S}_{\ell_3 n}[h] + \left(\mathcal{C}_{\ell_1 n}[f] \mathcal{C}_{\ell_2 n}[g] \mathcal{S}_{\ell_3 n}[h] + \text{cyclic} \right) \right], \quad (\text{A15})$$

where the sum runs over all 6 permutations of the three functions f, g, h . In this case, the total reduced bispectrum would be

$$b_{\ell_1 \ell_2 \ell_3} = \mathcal{A}_s^2 \sum_{n=0}^{n_{\max}} (\alpha_n \mathcal{C}_{\ell_1 \ell_2 \ell_3, n}[f, g, h] + \beta_n \mathcal{S}_{\ell_1 \ell_2 \ell_3, n}[f, g, h]). \quad (\text{A16})$$

If we have more terms with said structure, we can recover the full bispectrum by just summing them over.

At this point, one may wonder how much complication have we introduced with respect to the method presented in [42]. To see that, let's detail the expected computational sequence if we want to obtain the full bispectrum:

1. Compute the $\mathcal{C}_{\ell n}$ and $\mathcal{S}_{\ell n}$ for each ℓ and n we are interested in. In this extension, this must be done at worst three times per term, for three different f, g, h per term. That is at worst $3n_{\text{terms}}$ worst.
2. Further integrate on r the necessary combinations of the $\mathcal{C}_{\ell n}$ and $\mathcal{S}_{\ell n}$ to get the $\mathcal{C}_{\ell_1 \ell_2 \ell_3, n}$ and $\mathcal{S}_{\ell_1 \ell_2 \ell_3, n}$. In this case, this must be done once per term, times the 3 cyclic combinations of the functions, so $3n_{\text{terms}}$ again.
3. Decompose the shape functions $S(k_t)$ on Fourier modes and sum those over the $\mathcal{C}_{\ell_1 \ell_2 \ell_3, n}$ and $\mathcal{S}_{\ell_1 \ell_2 \ell_3, n}$. In the extension, this must be done once per term: n_{terms} slower.

Thus, given that the 2 first steps are the ones that take longest by far and dominate the computation time, the extension is $3n_{\text{terms}}$ as slow (or a smaller number of times n_{terms} if 2 or the three of f, g and h are the same, or if they are the same within a term of between terms).

However, if we are not interested in varying the background cosmology, and if the parameters of the primordial model enter through $S(k_t)$ only or as some external intensity factor f_{NL} , but not through the functions f, g and h , steps 1 and 2 can be pre-computed and stored. In that case, if one leaves the background cosmology unchanged, only step 3 is carried out and this method is only n_{terms} slower. Notice that the choice of the maximum order n_{\max} and the interval on which the Fourier decomposition is carried out depend implicitly on the model parameters; e.g. a higher frequency oscillation of the primordial shape requires a higher n_{\max} , which may not have been pre-computed yet.

In any case, remarkably, the computational costs grow only linearly with the number of terms, allowing us to compute more complicated non-separable bispectra that account for a richer set of physical scenarios.

3. Notes on precision

Every step in the computation described above, 1 to 3, has its own considerations regarding to precision: first the computation of the integrals of the $\mathcal{C}_{\ell n}$ and $\mathcal{S}_{\ell n}$, then the computation of $\mathcal{C}_{\ell_1 \ell_2 \ell_3, n}$ and $\mathcal{S}_{\ell_1 \ell_2 \ell_3, n}$, and finally, the Fourier decomposition of the shape functions.

a. *Computation of $\mathcal{C}_{\ell n}$ and $\mathcal{S}_{\ell n}$*

We want to compute with enough precision the integrals in eq. (A13). They have four elements: three oscillatory functions (a transfer function, a spherical Bessel function and a sine or cosine) and a coefficient function of a single k .

For starters, let's assume that the coefficient functions are very smooth compared to the rest of the oscillatory factors, so we can care about the sampling of the oscillations only. This is a reasonable hypothesis at least in the case that we are considering: they are monomials of a low degree.

For the oscillatory functions, we will parametrise the integration precision through an *adaptive* parameter N_I , meaning the number of samples per oscillation of the fastest oscillator for each combination of ℓ and n . In this paper, we use $N_I = 20$, which should be enough for most purposes. The integrals are performed using a Simpson integrator.

a. *Spherical Bessel functions:* They behave in the asymptotic limit as

$$\lim_{x \rightarrow 0} j_\ell(x) \sim x^\ell \quad \lim_{x \rightarrow \infty} j_\ell(x) \sim \frac{1}{x} \cos\left(x - (\ell + 1)\frac{\pi}{2}\right) \quad (\text{A17})$$

Therefore, to be sure to have N_I samples per oscillation, it is enough to have a $\delta x = \frac{2\pi}{N_I}$. In our particular case, the function is evaluated at $x := rk$, which means that the wavelength along k is $2\pi/r$. The distance along the line of sight, r is evaluated later over a small interval around recombination. The worst-case scenario, the shortest wavelength, corresponds to the maximum r of that interval. Thus, the desired sampling density on k will be

$$\delta k = \frac{2\pi}{N_I r_{\max}} \quad (\text{A18})$$

where r_{\max} is the maximum sampling value for the distance to recombination. For a reasonable value of $r_{\max} \simeq 1.5 \cdot 10^4$ Mpc, this means $\delta k \simeq 2 \cdot 10^{-5}$ Mpc $^{-1}$.

b. *Transfer functions:* They are roughly proportional to $j_\ell(r_{\text{re}}k)$, so the sampling strategy is the same as the last one, since r_{\max} is sampled very close to recombination, $r_{\max} \simeq r_{\text{re}}$.

c. *Fourier series basis:* Since the wavelength of the basis function of order n is $(k_{t,\max} - k_{t,\min})/n$, the necessary sampling here is

$$\delta k = \frac{k_{t,\max} - k_{t,\min}}{N_I n}. \quad (\text{A19})$$

Notice that for in interval of order 0 length we would need to go to an order $n \sim 10^4$ in the Fourier expansion for the basis functions to oscillate faster than the transfer and spherical Bessel functions. Thus, it will be the last ones that in most cases will impose the sampling density.

b. *Computation of $\mathcal{C}_{\ell_1 \ell_2 \ell_3, n}$ and $\mathcal{S}_{\ell_1 \ell_2 \ell_3, n}$*

The next step is to compute the integral along the line of sight from the recombination epoch to the present day in eq. (A14). Roughly speaking, the CMB temperature anisotropy is mainly produced by the inhomogeneities on the last-scattering surface, after photon-electron decoupling, so that the CMB photons are almost free to propagate until they reach us today. Mathematically, this means that each multipoles of the CMB radiation transfer functions behaves like a Dirac delta function centred at $\ell \sim r_{\text{rec}}k$. Hence, it is enough to sample a thin interval around $r_{\text{rec}} \simeq 1.4 \cdot 10^4$ Mpc [49]. In this work, we sample 200 points linearly spaced in the interval $[1.3304, 1.5284] \cdot 10^4$ Mpc, and integrate using a Simpson quadrature.

c. *Fourier decomposition: limits, order and sampling in ℓ*

a. *Limits of the Fourier decomposition:* In the choice of the limits of the interval of the Fourier decomposition, $[k_{t,\min}, k_{t,\max}]$, we must take into account two things. First, it is preferable that the shape functions are zero in both extremes of the interval; otherwise we can expect Gibbs overshoots and ringing in the extremes of the interval, and in order to suppress them we would need to go to higher order of the Fourier series, adding computational costs. Second, since the Fourier reconstruction is periodic outside the interval, if the interval is smaller than the sampling interval of the transfer functions, we may find copies of the shapes at higher k 's. Given that the maximum k_t reached is three times the maximum k_i sampled in the transfer functions, that have a sampling interval of roughly

$k_i \in [10^{-6}, 0.35] \text{ Mpc}^{-1}$, a good choice is to take the decomposition interval as $k_t \in [0, 1] \text{ Mpc}^{-1}$. As long as the Gibbs artefacts happen mostly at the end of the interval, they would be hidden by the low value of the transfer functions there.

b. Maximum order of the Fourier series: The order of the Fourier series must be high enough to correctly represent the shape functions, which is completely model dependent. This is best checked by directly comparing the original bispectrum shape with the reconstruction from the Fourier decomposition.

c. Sampling in ℓ space: Approximately, $\ell \simeq r_{\text{re}} k$; thus, if a sampling density δk accurately represents the primordial bispectrum, a corresponding $\delta \ell = r_{\text{re}} \delta k$, should provide with a good sample of the bispectrum. For computational feasibility, we may reduce the sampling density in ℓ and construct a sufficiently accurate spline approximation, in order to calculate e.g. the signal-to-noise.

4. Application to our case

Let us repeat the equation of the primordial bispectrum:

$$B_{\mathcal{R}}(k_1, k_2, k_3) = |u|_{\text{max}} \frac{(2\pi)^4 \mathcal{A}_s^2}{(k_1 k_2 k_3)^2} \sum_{i=0}^2 c_i(k_1, k_2, k_3) \left(\frac{d}{d^{k_t/2}} \right)^i \frac{1}{|u|_{\text{max}}} \frac{\Delta \mathcal{P}_{\mathcal{R}}}{\mathcal{P}_{\mathcal{R}}}(k_t/2; |s|_{\text{max}}, \tau_0), \quad (\text{A20})$$

where $k_t := k_1 + k_2 + k_3$. The c_i coefficients are given in equation (10).

We have made explicit the dependency on the parameters of the reduction in c_s , $(|u|_{\text{max}}, |s|_{\text{max}}, \tau_0)$, to highlight an important property of this expression:

- $|u|_{\text{max}}$ enters only as an overall factor (the combination $|u|_{\text{max}}^{-1} \frac{\Delta \mathcal{P}_{\mathcal{R}}}{\mathcal{P}_{\mathcal{R}}}$ does not depend on it).
- $|s|_{\text{max}}$ and τ_0 enter only through $|u|_{\text{max}}^{-1} \frac{\Delta \mathcal{P}_{\mathcal{R}}}{\mathcal{P}_{\mathcal{R}}}$.

The rest of the factors do not depend on the particular choice of the reduction. Thus, we can pre-compute all the integrals $\mathcal{C}_{\ell_1 \ell_2 \ell_3, n}$ and $\mathcal{S}_{\ell_1 \ell_2 \ell_3, n}$ for all function combinations and compute the Fourier decomposition and sum the terms on the fly.

We also need to care about two model-dependent aspects of the precision: the maximum order of the Fourier decomposition n_{max} and the necessary sampling in ℓ (so we can interpolate for a faster computation of S/N). The determinant quantity in both cases is the maximum oscillatory frequency of the feature, and thus the maximum $|\tau_0|$ that we are interested in.

For n_{max} , notice that the *characteristic* order of the feature, i.e. that whose Fourier frequency is equals the oscillatory frequency of the feature, is $|\tau_0|(k_{t, \text{max}} - k_{t, \text{min}})(2\pi)^{-1}$. Decomposing the shape functions after twice that frequency should provide us with a good reconstruction. For the sake of safety, we may increase that order by a small security factor:

$$n_{\text{max}} = |\tau_0|(k_{t, \text{max}} - k_{t, \text{min}}) \frac{1}{\pi} (1 + \epsilon) \quad (\text{A21})$$

For the maximum $|\tau_0|$ we are interested in in this paper, 1600, an interval of 1 Mpc^{-1} , as stated in the previous section, and a security factor of 1.5, the maximum order that we need to pre-compute is $n_{\text{max}} = 765$ (400 is enough for the high β regime).

Regarding the sample in the ℓ -space, notice first that, approximately, $\ell \simeq r_{\text{re}} k$. In the ℓ -space, we aim at correctly sampling the feature oscillation. With the relationship stated above, the wavelength in ℓ of the feature is, approximately, $r_{\text{re}} 2\pi/|\tau_0|$. This means that if we wanted a reasonable amount of 20 samples per feature oscillation, to get a good interpolation, we would need a sampling $\Delta \ell \simeq 3$.

Here, we may notice that we can take advantage of the two different regimes discussed in section REFERENCE: the features of highest β have also a lower $|\tau_0|$, maximum around 820, which means that $\Delta \ell \simeq 5$ is enough sampling. On the other hand, the features at lowest β , though they need a sampling of $\Delta \ell \simeq 3$ and higher, are dead by $k_t = 0.03 \text{ Mpc}^{-1}$, which corresponds roughly to $\ell_t \simeq 400$. Thus we can sample more finely up to that ℓ_t , and more coarsely after it, allowing for a higher computational efficiency.

5. Fisher matrix elements, signal-to-noise and correlations

We will use the spectra predicted/obtained to compute signal-to-noise and shape correlations through the Fisher matrix [49] between two temperature bispectra i and j , assuming homogeneous noise:

$$F_{ij} := \sum_{\ell_{\min} \leq \ell_1 \leq \ell_2 \leq \ell_3 \leq \ell_{\max}} \frac{2}{\pi} \left(\ell_1 + \frac{1}{2} \right) \left(\ell_2 + \frac{1}{2} \right) \left(\ell_3 + \frac{1}{2} \right) \begin{pmatrix} \ell_1 & \ell_2 & \ell_3 \\ 0 & 0 & 0 \end{pmatrix}^2 \frac{b_{\ell_1 \ell_2 \ell_3}^{(i)} b_{\ell_1 \ell_2 \ell_3}^{(j)}}{\sigma_{\ell_1 \ell_2 \ell_3}^2}, \quad (\text{A22})$$

where ℓ_{\min} and ℓ_{\max} are the minimum and maximum values allowed for individual multipoles, here respectively 2 and 2000, and $\sigma_{\ell_1 \ell_2 \ell_3}^2$ is the approximate cosmic variance for a small bispectrum [50–53]

$$\sigma_{\ell_1 \ell_2 \ell_3}^2 \approx C_{\ell_1} C_{\ell_2} C_{\ell_3} \Delta_{\ell_1 \ell_2 \ell_3} / f_{\text{sky}}, \quad (\text{A23})$$

where C_{ℓ_i} is the *observed* spectrum $C_{\ell, \text{obs}} = C_{\ell, \text{theo}} + N_{\ell} b_{\ell}^{-2}$, which we have assumed an effective Gaussian beam with $\theta_{\text{FWHM}} = 7.25$ arcmin and white noise with standard deviation $\sigma_N = 33 \mu\text{K arcmin}$, and the map mask leaves a sky fraction $f_{\text{sky}} = 0.76$. The geometrical factor $\Delta_{\ell_1 \ell_2 \ell_3}$ evaluates to 1, 2, 6 respectively for the cases of all ℓ 's being different, two being equal and all being equal. For i and j being two different bispectra, we assume the same background cosmology, and only different inflationary model. We compute the necessary Wigner 3- j symbol using the WIGXJPF algorithm [54].

From this Fisher matrix, one can derive the signal-to-noise and the correlation between two bispectra i and j as [55]

$$\left(\frac{\text{S}}{\text{N}} \right)_i = \sqrt{F_{i,i}} \quad \text{and} \quad C_{i,j} = \frac{F_{i,j}}{\sqrt{F_{i,i} F_{j,j}}}. \quad (\text{A24})$$

Notice that the correlation between two bispectra is independent of the amplitude of either.

We compute the Fisher matrix elements summing by slices of constant $\ell_t = \ell_1 + \ell_2 + \ell_3$, with $\ell_i \in [2, 2000]$, and interpolating the values we have not sampled in our bispectrum computation. We sample as many slices as we can within the target error for the Fisher matrix elements (5–10%).

In principle, we could have taken advantage of the actual separability of the modal-expanded shape to pre-compute some of the steps of the Fisher matrix computation, and even maybe to create a KSW estimator, as it is done in [42]. However, we choose not to do so in the present work: the amount of pre-computation needed is high, due to the sizeable number of different combinations of c_i coefficients, and for now we do not want to streamline the Fisher matrix element computations – we are not scanning the parameter space of the feature, but just computing signal-to-noise and correlations for particular parameters combinations.

-
- [1] R. Adam et al. (Planck), *Astron. Astrophys.* **594**, A1 (2016), 1502.01582.
 - [2] P. A. R. Ade et al. (Planck) (2015), 1502.02114.
 - [3] P. A. R. Ade et al. (Planck) (2015), 1502.01592.
 - [4] P. A. R. Ade et al. (Planck), *Astron. Astrophys.* **571**, A22 (2014), 1303.5082.
 - [5] P. A. R. Ade et al. (Planck), *Astron. Astrophys.* **571**, A24 (2014), 1303.5084.
 - [6] J. R. Fergusson, M. Liguori, and E. P. S. Shellard, *JCAP* **12**, 032 (2012), 1006.1642.
 - [7] X. Chen, R. Easther, and E. A. Lim, *JCAP* **0706**, 023 (2007), astro-ph/0611645.
 - [8] J. Chluba, J. Hamann, and S. P. Patil, *Int. J. Mod. Phys.* **D24**, 1530023 (2015), 1505.01834.
 - [9] A. Achúcarro, J.-O. Gong, G. A. Palma, and S. P. Patil, *Phys. Rev.* **D87**, 121301 (2013), 1211.5619.
 - [10] J.-O. Gong, K. Schalm, and G. Shiu, *Phys. Rev.* **D89**, 063540 (2014), 1401.4402.
 - [11] G. A. Palma, *JCAP* **1504**, 035 (2015), 1412.5615.
 - [12] S. Mooij, G. A. Palma, G. Panotopoulos, and A. Soto, *JCAP* **1510**, 062 (2015), [Erratum: *JCAP*1602,no.02,E01(2016)], 1507.08481.
 - [13] J. R. Fergusson, H. F. Gruetjen, E. P. S. Shellard, and M. Liguori, *Phys. Rev.* **D91**, 023502 (2015), 1410.5114.
 - [14] J. R. Fergusson, H. F. Gruetjen, E. P. S. Shellard, and B. Wallisch, *Phys. Rev.* **D91**, 123506 (2015), 1412.6152.
 - [15] P. D. Meerburg, M. Münchmeyer, and B. Wandelt, *Phys. Rev.* **D93**, 043536 (2016), 1510.01756.
 - [16] S. Appleby, J.-O. Gong, D. K. Hazra, A. Shafieloo, and S. Sypsas, *Phys. Lett.* **B760**, 297 (2016), 1512.08977.
 - [17] A. Achúcarro, V. Atal, P. Ortiz, and J. Torrado, *Phys. Rev.* **D89**, 103006 (2014), 1311.2552.
 - [18] A. Achúcarro, V. Atal, B. Hu, P. Ortiz, and J. Torrado, *Phys. Rev.* **D90**, 023511 (2014), 1404.7522.
 - [19] B. Hu and J. Torrado, *Phys. Rev.* **D91**, 064039 (2015), 1410.4804.
 - [20] C. Cheung, P. Creminelli, A. L. Fitzpatrick, J. Kaplan, and L. Senatore, *JHEP* **03**, 014 (2008), 0709.0293.

- [21] A. Achúcarro, J.-O. Gong, S. Hardeman, G. A. Palma, and S. P. Patil, Phys. Rev. **D84**, 043502 (2011), 1005.3848.
- [22] S. Cespedes, V. Atal, and G. A. Palma, JCAP **1205**, 008 (2012), 1201.4848.
- [23] A. Achúcarro, V. Atal, S. Cespedes, J.-O. Gong, G. A. Palma, and S. P. Patil, Phys. Rev. **D86**, 121301 (2012), 1205.0710.
- [24] D. Seery and J. E. Lidsey, JCAP **0506**, 003 (2005), astro-ph/0503692.
- [25] S. Weinberg, Phys. Rev. **D72**, 043514 (2005), hep-th/0506236.
- [26] J. M. Maldacena, JHEP **05**, 013 (2003), astro-ph/0210603.
- [27] P. Adshead and W. Hu, Phys. Rev. **D89**, 083531 (2014), 1402.1677.
- [28] D. Cannone, N. Bartolo, and S. Matarrese, Phys. Rev. **D89**, 127301 (2014), 1402.2258.
- [29] V. Miranda and W. Hu, Phys. Rev. **D89**, 083529 (2014), 1312.0946.
- [30] P. Creminelli and M. Zaldarriaga, JCAP **0410**, 006 (2004), astro-ph/0407059.
- [31] A. Lewis, A. Challinor, and A. Lasenby, Astrophys. J. **538**, 473 (2000), astro-ph/9911177.
- [32] C. Howlett, A. Lewis, A. Hall, and A. Challinor, JCAP **1204**, 027 (2012), 1201.3654.
- [33] N. Aghanim et al. (Planck), Astron. Astrophys. **594**, A11 (2016), 1507.02704.
- [34] W. J. Handley, M. P. Hobson, and A. N. Lasenby, Mon. Not. Roy. Astron. Soc. **453**, 4384 (2015), 1506.00171.
- [35] A. Lewis and S. Bridle, Phys. Rev. **D66**, 103511 (2002), astro-ph/0205436.
- [36] F. Feroz and M. P. Hobson, Mon. Not. Roy. Astron. Soc. **384**, 449 (2008), 0704.3704.
- [37] F. Feroz, M. P. Hobson, and M. Bridges, Mon. Not. Roy. Astron. Soc. **398**, 1601 (2009), 0809.3437.
- [38] F. Feroz, M. P. Hobson, E. Cameron, and A. N. Pettitt (2013), 1306.2144.
- [39] P. A. R. Ade et al. (Planck), Astron. Astrophys. **594**, A13 (2016), 1502.01589.
- [40] Y.-F. Cai, E. G. M. Ferreira, B. Hu, and J. Quintin, Phys. Rev. **D92**, 121303 (2015), 1507.05619.
- [41] M. Benetti and J. S. Alcaniz, Phys. Rev. **D94**, 023526 (2016), 1604.08156.
- [42] M. Münchmeyer, P. D. Meerburg, and B. D. Wandelt, Phys. Rev. **D91**, 043534 (2015), 1412.3461.
- [43] M. Liguori, E. Sefusatti, J. R. Fergusson, and E. P. S. Shellard, Adv. Astron. **2010**, 980523 (2010), 1001.4707.
- [44] J. R. Fergusson, M. Liguori, and E. P. S. Shellard, JCAP **1212**, 032 (2012), 1006.1642.
- [45] X. Chen, P. D. Meerburg, and M. Münchmeyer, JCAP **1609**, 023 (2016), 1605.09364.
- [46] Y. Xu, J. Hamann, and X. Chen (2016), 1607.00817.
- [47] X. Chen, C. Dvorkin, Z. Huang, M. H. Namjoo, and L. Verde, JCAP **1611**, 014 (2016), 1605.09365.
- [48] M. Ballardini, F. Finelli, C. Fedeli, and L. Moscardini, JCAP **1610**, 041 (2016), 1606.03747.
- [49] K. M. Smith and M. Zaldarriaga, Mon. Not. Roy. Astron. Soc. **417**, 2 (2011), astro-ph/0612571.
- [50] X.-c. Luo, Astrophys. J. **427**, L71 (1994), astro-ph/9312004.
- [51] A. F. Heavens, Mon. Not. Roy. Astron. Soc. **299**, 805 (1998), astro-ph/9804222.
- [52] D. N. Spergel and D. M. Goldberg, Phys. Rev. **D59**, 103001 (1999), astro-ph/9811252.
- [53] A. Gangui and J. Martin, Mon. Not. Roy. Astron. Soc. **313**, 323 (2000), astro-ph/9908009.
- [54] H. T. Johansson and C. Forssén, SIAM J. Sci. Statist. Comput. **38**, A376 (2016), 1504.08329.
- [55] E. Komatsu and D. N. Spergel, Phys. Rev. **D63**, 063002 (2001), astro-ph/0005036.